

Designing a Campus

This chapter covers design considerations and recommendations for implementing the Cisco AVVID network in a campus environment.

Campus Switching Designs for Cisco AVVID

Until recently, the conventional wisdom was that Quality of Service (QoS) would never be an issue in the enterprise campus due to the bursty nature of network traffic and the capability of buffer overflow. Gradually, engineers have come to understand that buffering, not bandwidth, is the issue in the campus. For this reason, QoS tools are required to manage these buffers to minimize loss, delay, and delay variation. Figure 3-1 shows areas where transmit buffers can give rise to QoS issues.

Figure 3-1 QoS Considerations with Transmit Buffers

Transmit buffers have a tendency to fill to capacity in high-speed campus networks due to the bursty nature of data networks combining with the high volume of smaller Transmission Control Protocol (TCP) packets. If an output buffer fills, ingress interfaces are not able to place new flow traffic into the output buffer. Once the ingress buffer fills, which can happen very quickly, packet drops will occur. Typically, these drops are more than a single packet in any given flow. As stated earlier, packet loss causes voice clipping and skips. Current Cisco Digital Signal Processor (DSP) algorithms can correct for 30 ms of lost voice. Cisco VoIP technology uses 20-ms samples of voice payload per VoIP packet. Thus, current DSP algorithms allow for only a single voice Real-time Transport Protocol (RTP) packet to be lost during any given time. If two successive voice packets are lost, voice quality begins to degrade. Figure 3-2 illustrates this situation.

Figure 3-2 *Loss of Voice Quality When Transmit Buffer Is Full*

VoIP traffic is sensitive to both delayed packets and dropped packets. As long as a campus is using Gigabit Ethernet trunks, which have extremely fast serialization times, delay should never be a factor regardless of the size of the queue buffer. Drops, however, always adversely affect voice quality in the campus. Using multiple queues on transmit interfaces is the only way to eliminate the potential for dropped traffic caused by buffers operating at 100% capacity. By separating voice and video (which are both sensitive to delays and drops) into their own queues, you can prevent flows from being dropped at the ingress interface even if data flows are filling up the data transmit buffer. Figure 3-3 illustrates the use of separate voice and data buffers.

Figure 3-3 *Using Separate Transmit Buffers for Voice and Data*



Note

It is critical to verify that Flow Control is disabled when enabling QoS (multiple queues) on Catalyst switches. Flow Control will interfere with the configured queuing behavior by acting on the ports before queuing is activated. Flow Control is disabled by default.

Queue Scheduling

The scheduler process can use a variety of methods to service each of the transmit queues (voice and data). The easiest method is a Round-Robin (RR) algorithm, which services queue 1 through queue N in a sequential manner. While not robust, this is an extremely simple and efficient method that can be used for branch office and wiring closet switches. Distribution Layer switches use a Weighted Round-Robin (WRR) algorithm in which higher priority traffic is given a scheduling "weight."

Another option is to combine Round-Robin or Weighted Round-Robin scheduling with priority scheduling for applications that are sensitive to packet delay and drop. This uses a priority queue (PQ) that is always served first when there are packets in the queue. If there are no frames in the PQ, the additional queues are scheduled using RR or WRR.

Number of Queues

There has been much discussion about how many queues are actually needed on transmit interfaces in the campus. Should you add a queue to the wiring closet switches for each Class of Service (CoS) value? Should you add eight queues to the distribution layer switches? Should you add a queue for each of the 64 Differentiated Services Code Point (DSCP) values? This section presents some guidelines that address these questions.

First, it is important to remember that each port has a finite amount of buffer memory. A single queue has access to all the memory addresses in the buffer. As soon as a second queue is added, the finite buffer amount is split into two portions, one for each queue. Now all packets entering the switch must contend for a much smaller portion of buffer memory. During periods of high traffic, the buffer fills, and packets are dropped at the ingress interface. Because the majority of network traffic today is TCP-based, a dropped packet results in a re-send, which further increases network congestion. Therefore, queuing should be used cautiously and only when particular priority traffic is sensitive to packet delays and drops.

Two queues are adequate for wiring closet switches, where buffer management is less critical than at other layers. How these queues are serviced (Round-Robin, Weighted Round-Robin, or Priority Queuing) is less critical than the number of buffers because the scheduler process is extremely fast when compared to the aggregate amount of traffic.

Distribution layer switches require much more complex buffer management due to the flow aggregation occurring at that layer. Not only are priority queues needed, but you should also specify thresholds within the standard queues. Cisco has chosen to use multiple thresholds within queues instead of continually increasing the number of interface queues. As discussed earlier, each time a queue is configured and allocated, all of the memory buffers associated with that queue can be used only by frames meeting the queue entrance criteria. The following example illustrates this concept:

Assume that a Catalyst 4000 10/100 Ethernet port has two queues configured, one for VoIP (VoIP bearer and control traffic) and the default queue, which is used for Hypertext Transfer Protocol (HTTP), e-mail, File Transfer Protocol (FTP), logins, Windows NT Shares, and Network File System (NFS). The 128-KB voice queue is split into a 7:1 transmit and receive ratio. The transmit buffer memory is then further separated into high- and low-priority partitions in a 4:1 ratio. If the default traffic (the web, e-mail, and file shares) begins to congest the default queue, which is only 24 KB, packets begin dropping at the ingress interfaces. This occurs

regardless of whether the VoIP control traffic is using any of its queue buffers. The dropped packets of the TCP-oriented applications cause each of these applications to send the data again, aggravating the congested condition of the network. If this same scenario were configured with a single queue, but with multiple thresholds used for congestion avoidance, the default traffic would share the entire buffer space with the VoIP control traffic. Only during periods of congestion, when the entire buffer memory approaches saturation, would the lower priority traffic (HTTP and e-mail) be dropped.

This discussion does not imply that multiple queues are to be avoided in Cisco AVVID networks. As discussed earlier, the VoIP bearer streams *must* use a separate queue to eliminate the adverse affects that packet drops and delays have on voice quality. However, every single CoS or DSCP value should *not* get its own queue because the small size of the resulting default queue will cause many TCP re-sends and will actually increase network congestion.

In addition, the VoIP bearer channel is a bad candidate for queue congestion avoidance algorithms such as Weighted Random Early Detection (WRED). Queue thresholding uses the WRED algorithm to manage queue congestion when a preset threshold value is specified. Random Early Detection (RED) works by monitoring buffer congestion and discarding TCP packets if the congestion begins to increase. The result of the drop is that the sending endpoint detects the dropped traffic and slows the TCP sending rate by adjusting the window size. A WRED drop threshold is the percentage of buffer utilization at which traffic with a specified CoS value is dropped, leaving the buffer available for traffic with higher priority CoS values. The key is the word "Random" in the algorithm name. Even with weighting configured, WRED can still discard packets in any flow; it is just statistically more likely to drop them from the lower CoS thresholds.

Marking Control and Management Traffic

In networks with high traffic loads, managing the delivery of control traffic is critical to ensuring a positive user experience with VoIP. An example where this comes into play is with the Delay to Dial-Tone (DTT) time. The Cisco IP Phones use Skinny Station Protocol to communicate with Cisco CallManager. When a Cisco IP Phone goes off hook, it "asks" Cisco CallManager what to do. Cisco CallManager then instructs the Cisco IP Phone to play dial-tone. If this Skinny Client Protocol management and control traffic is dropped or delayed within the network, the user experience is adversely affected. This same logic applies to all signaling traffic for gateways and phones.

To ensure that this control and management traffic is marked as important (but not as important as the actual RTP stream), Access Control Lists (ACLs) are used to classify these streams on Catalyst 5000 and 6000 switches that are enabled for Layer 3/4. Examples of these configurations are detailed in the “Catalyst 6000 Access Layer” section on page 3-11. For designs where a Cisco IOS router is the first Layer 3 or 4 access point, ACLs are used. Examples of these configurations are included in Chapter 5, “Implementing a Wide Area Network.”

Figure 3-4 shows a typical server farm design, with an access layer switch providing access to the distribution layer.

Figure 3-4 *Typical Server Farm*

Skinny Protocol

Cisco CallManager communicates with IP phones and gateways using TCP ports 2000-2002. The following example commands classify all Skinny Protocol traffic from IP phones and gateways (VLAN 110) and Cisco CallManager (4/2) as DSCP 26 (AF31, which is backward compatible with IP Precedence 3).



Note

Beginning with Release 3.0(5), Cisco CallManager includes the ability to configure the CoS and ToS values for all VoIP control and management traffic from Cisco CallManager, the IP phones, and the Skinny Protocol gateways (this does not include the AT and AS model analog gateways). With this user-configurable classification, network element access lists are no longer required for marking Skinny Protocol VoIP control traffic. H.323 and Media Gateway Control Protocol (MGCP) traffic still require external, network element marking for several more months.

The following commands perform these functions:

1. Enable switch-wide QoS.
2. Create an access control list (ACL_IP-PHONES), marking all Skinny Client and Gateway Protocol traffic from the IP phones and from Skinny Protocol gateways with a DSCP value of 26 (AF31).
3. Add to the ACL_IP-PHONE access list, trusting all DSCP markings from the IP phone, so that the ToS=5 RTP traffic is not rewritten.
4. Create an access control list (ACL_VOIP_CONTROL), marking all Skinny Client and Gateway Protocol traffic from Cisco CallManager with a DSCP value of 26 (AF31).
5. Accept incoming Layer 2 CoS classification. (Current 10/100 version "1" line cards must have **trust-cos** enabled even though the parser returns an error).
6. Inform the port that all QoS associated with the port will be done on a VLAN basis.
7. Instruct the IP phone to rewrite CoS from the PC to CoS=0 within the IP phone Ethernet ASIC.
8. Inform Cisco CallManager port (4/2) that all QoS associated with the port will be done on a port basis.

9. Write the access control list to hardware.
10. Map the ACL_IP-PHONE access control list to the auxiliary VLAN.
11. Map the ACL_VOIP_CONTROL access control list to the Cisco CallManager port.

```
cat6k-access> (enable) set qos enable
cat6k-access> (enable) set qos acl ip ACL_IP-PHONES dscp 26 tcp any any range 2000 2002
cat6k-access> (enable) set qos acl ip ACL_IP-PHONES trust-cos ip any any
cat6k-access> (enable) set qos acl ip ACL_VOIP_CONTROL dscp 26 tcp any any range 2000
2002
cat6k-access> (enable) set port qos 5/1-48 trust trust-cos
cat6k-access> (enable) set port qos 5/1-48 vlan-based
cat6k-access> (enable) set port qos 5/1-48 trust-ext untrusted
cat6k-access> (enable) set port qos 4/2 port-based
cat6k-access> (enable) commit qos acl all
cat6k-access> (enable) set qos acl map ACL_IP-PHONES 110
cat6k-access> (enable) set qos acl map ACL_VOIP_CONTROL 4/2
```

H.323 Protocol

Cisco CallManager communicates with H.323 gateways using TCP ports 1720 (H.225) and 11xxx (H.245). The following example commands classify H.323 control traffic from Cisco CallManager (4/2) and from H.323 gateways (4/3) as DSCP 26 (AF31, which is backward compatible with IP Precedence 3).

```
cat6k-access> (enable) set qos acl ip ACL_VOIP_CONTROL dscp 26 tcp any any eq 1720
cat6k-access> (enable) set qos acl ip ACL_VOIP_CONTROL dscp 26 tcp any any range 11000
11999
cat6k-access> (enable) set port qos 4/2 port-based
cat6k-access> (enable) set port qos 4/3 port-based
cat6k-access> (enable) commit qos acl ACL_VOIP_CONTROL
cat6k-access> (enable) set qos acl map ACL_VOIP_CONTROL 4/2
cat6k-access> (enable) set qos acl map ACL_VOIP_CONTROL 4/3
```

MGCP

Cisco CallManager communicates with Media Gateway Control Protocol (MGCP) gateways using User Datagram Protocol (UDP) port 2427. The following example commands classify MGCP control traffic from Cisco CallManager (4/2) and from the MGCP gateway (4/4) as DSCP 26 (AF31, which is backward compatible with IP Precedence 3).

```
cat6k-access> (enable) set qos acl ip ACL_VOIP_CONTROL dscp 26 udp any any eq 2427
cat6k-access> (enable) set port qos 4/2 port-based
cat6k-access> (enable) set port qos 4/4 port-based
cat6k-access> (enable) commit qos acl ACL_VOIP_CONTROL
cat6k-access> (enable) set qos acl map ACL_VOIP_CONTROL 4/2
cat6k-access> (enable) set qos acl map ACL_VOIP_CONTROL 4/4
```

Example 3-1 shows the command and its associated output for verifying that the Access Control Lists (ACLs) are attached to the correct VLANs and ports.

Example 3-1 Verifying the ACLs

```
cat6k-access> (enable) sh qos acl map run all
```

ACL name	Type	Vlans
-----	-----	-----
ACL_IP-PHONES	IP	110,111,112
ACL name	Type	Ports
-----	-----	-----
ACL_IP-PHONES	IP	
ACL name	Type	Vlans
-----	-----	-----
ACL_VOIP_CONTROL	IP	
ACL name	Type	Ports
-----	-----	-----
ACL_VOIP_CONTROL	IP	4/2,4/3,4/4

Catalyst 6000 Access Layer

One of the most popular campus configurations for Cisco AVVID solutions is to use Catalyst 6000 switches in both the wiring closet and the distribution and core layers. There are several compelling reasons for this:

- The Catalyst 6000 can provide in-line power to the IP phones.
- The Catalyst 6000 offers the highest growth potential.
- The Catalyst 6000 supports the most advanced Layer 2/3 campus QoS tools in the Cisco product line.

Figure 3-5 shows a general model for the Catalyst 6000 QoS configurations discussed in this guide.

Figure 3-5 *General Model for Catalyst 6000 QoS Configurations*

With the addition of the Policy Feature Card (PFC) daughter card, the Catalyst 6000 is inherently capable of handling Layer 2, 3, and 4 QoS issues. The PFC can be used to enable advanced QoS tools such as packet classification and marking, scheduling, and congestion avoidance based on either Layer 2 or Layer 3 and 4 header information. Multiple receive and transmit queues with thresholds can be configured and used according to the QoS policy rules configured in the switch.

The Catalyst 6000 has two versions of Supervisor Engines, the Sup1 and Sup1A. There are also two versions of Catalyst 6000 line cards, the second of which is also denoted by an "A" product number. All Catalyst 6000 Ethernet modules support a single receive queue with four thresholds and two transmit queues, each with two thresholds. The "A" cards include enhanced QoS features that provide an additional priority queue for both ingress and egress interfaces. These queues are serviced in a Weighted Round-Robin (WRR) method, except for the priority queue, which is always serviced as soon as frames have entered the queue. To see how a port is configured, issue the **show port capabilities <mod/port>** CatOS command. The default QoS capabilities of the port can be changed using the **set qos map <port_type> rx | tx <queue#> <threshold#>** and **set qos wred-threshold** commands. When modifying the queue thresholds, it is important to remember that the higher priority queue has a higher numerical value.

Scheduling for the Catalyst 6000 transmit interfaces is managed by the WRR algorithm. Each queue is given a user-configurable weight. By default, the "high" queue is given 98% of the scheduler time, and the "low" queue is given just 2%. This ratio is conducive to ensuring that packets with a low delay tolerance are not delayed in a queue. This is also the reason behind giving the "low" queue a much higher percentage of the overall interface buffer. If the Priority Queue (PQ) is configured, it will always be serviced first. If no frames reside in the PQ, WRR begins to schedule the other two queues.

Catalyst 6000 Port Scheduling and Queuing Schemes

This section presents several suggested configurations for port scheduling and queuing on the Catalyst 6000.

Receive Interface

There are two possible configurations for the receive interface, depending on whether or not a priority queue is needed:

- 1Q4T

One standard queue with four drop thresholds.

8-KB receive buffer for 10/100 Mbps

64-KB receive buffer for 1000 Mbps

Available on all 10/100/1000 Mbps modules

The default values for the drop thresholds are

% of Buffer Capacity	Drop CoS Value
50%	0-1
60%	2-3
80%	4-5
100%	6-7

- 1P1Q4T

One Priority Queue (PQ) and one standard queue with four drop thresholds

Available only on certain versions of 10/100/1000 Mbps modules, depending on line card

By default, all CoS 5 frames are placed in the PQ, which uses a strict priority scheduling algorithm.

The default values for the drop thresholds in the standard queue are

Queue #	% of Buffer Capacity	Drop CoS Value
1	50%	0-1
1	60%	2-3
1	80%	4
1	100%	6-7
2	100%	5

Transmit Interface

There are two possible configurations for the transmit interface, depending on whether a priority queue is needed:

- 2Q2T

Two standard queues with two drop thresholds. The high-priority queue is allocated 20% of the total queue size, and the low-priority queue is allocated 80% of the total queue size.

Available on all 10/100/1000 Mbps modules

The default values for the drop thresholds are

Queue #	% of Buffer Capacity	Drop CoS Value
1 - Low Priority - 80% of total queue size	40%	0-1
	100%	2-3
2 - High Priority - 20% of total queue size	40%	4-5
	100%	6-7

- 1P2Q2T

One Priority Queue (PQ) and two standard queues with two drop thresholds. By default, all CoS 5 frames are placed in the PQ, which uses a strict priority scheduling algorithm that always services the PQ first, and, once the PQ is

empty, WRR is used on the remaining queues. The PQ gets allocated 15% of the total queue size, as does the high-priority queue. The low-priority queue is allocated 70% of the total queue size.

Available only on certain versions of 10/100/1000 Mbps modules, depending on line card

The default values for the drop thresholds are

Queue #	% of Buffer Capacity	Drop CoS Value
1 - Low Priority - 70% of total queue size	40%	0-1
	100%	2-3
2 - High Priority - 15% of total queue size	40%	4
	100%	6-7
3 - Priority Queue - 15% of total queue size	100%	5

Configuring QoS Parameters

After you have connected the IP phone to the wiring closet switch (see Chapter 2, “Connecting IP Phones”), it is time to configure the QoS parameters on the switch. This includes setting up multiple queues on all ports, configuring access to the queues, setting thresholds for traffic drops, and connecting the switch to the distribution or core layer. The following sections detail these steps.

IP Phone Port Queuing

If you use a single cable to connect IP phone, as illustrated in Figure 3-6, the access port is configured to trust the IP phone and not the attached PC. The port is also configured to use multiple transmit queues, one being a priority queue for voice traffic.

Figure 3-6 Using a Single Cable to Connect an IP Phone

The following commands enable QoS on the access layer Catalyst 6000 by performing these functions:

1. Enable switch-wide QoS.
2. Inform the port that all QoS associated with the port will be done on a VLAN basis.
3. Instruct the IP phone to rewrite CoS from the PC to CoS=0 within the IP phone Ethernet ASIC.
4. Accept incoming Layer 2 CoS classification. (Current 10/100 version "1" line cards must still have **trust-cos** enabled even though the parser returns an error).
5. Create an access list that accepts incoming Layer 3 ToS classification (necessary only on 10/100 ports).
6. Write the access list to hardware.
7. Map the access list to the auxiliary VLAN.

```
cat6k-access> (enable) set qos enable
cat6k-access> (enable) set port qos 5/1-48 vlan-based
cat6k-access> (enable) set port qos 5/1-48 trust-ext untrusted
cat6k-access> (enable) set port qos 5/1-48 trust trust-cos
cat6k-access> (enable) set qos acl ip ACL_IP-PHONES trust-cos any
cat6k-access> (enable) commit qos acl ACL_IP-PHONES
cat6k-access> (enable) set qos acl map ACL_IP-PHONES 110
```

Once QoS has been enabled on the Catalyst 6000 access layer switch, you can use the following command to place all CoS=3 (VoIP control) traffic into the second transmit queue, with a low drop threshold, to ensure successful call control during periods of heavy congestion. All CoS=5 (VoIP RTP Bearer) traffic is placed into the second queue automatically.

```
cat6k-access> (enable) set qos map 2q2t tx 2 1 cos 3
```

Verifying IP Phone Access Port Configuration

One of the fundamental processes of implementing Quality of Service (QoS) is verifying that the configurations are actually performing as expected. On the Catalyst 6000 access layer switch, you can verify configuration performance during periods of high congestion by examining the output of the following commands:

- `show port qos <mod/port>`

This command shows the QoS settings for the specified port. See Example 3-2.

- `show qos info runtime <mod/port>`

This command shows QoS runtime information for the specified port. See Example 3-3.

- `show mac <mod/port>`

This command shows Media Access Control (MAC) information for the specified port. See Example 3-4.

- `show qos statistics l3`

This command shows summary QoS statistics for all ports. See Example 3-5.

- `show qos stat <mod/port>`

This command shows detailed QoS statistics for the specified port. See Example 3-6.

Example 3-2 Displaying QoS Settings

```
cat6k-access> (enable) sh port qos 5/1
QoS is enabled for the switch
QoS policy source for the switch set to local.
```

Port	Interface config	Type	Interface runtime	Type	Policy config	Source	Policy runtime	Source
5/1	vlan-based		vlan-based			COPS		local

Port	TxPort	Type	RxPort	Type	Trust config	Type	Trust runtime	Type	Def CoS config	Def CoS runtime
5/1	2q2t		1q4t		trust-cos		trust-cos*		0	0

Port	Ext-Trust	Ext-Cos
5/1	untrusted	0

(*)Runtime trust type set to untrusted.

Config:

Port	ACL name	Type
No ACL is mapped to port 5/1.		
ACL is mapped to VLAN		

Runtime:

Port	ACL name	Type
No ACL is mapped to port 5/1.		

Example 3-3 Displaying QoS Runtime Information

```

cat6k-access>(enable) sh qos info run 5/1
Run time setting of QoS:
QoS is enabled
Policy Source of port 5/1: Local
Current 10/100 "1" linecards support 2q2t/1q4t only
Tx port type of port 5/1 : 2q2t
Rx port type of port 5/1 : 1q4t
Interface type: vlan-based
ACL is mapped to VLAN
ACL attached:
The qos trust type is set to trust-cos.
Warning: Runtime trust type set to untrusted.
Default CoS = 0
Queue and Threshold Mapping for 2q2t (tx):
Queue Threshold CoS
-----
1      1      0 1
1      2      2
2      1      3 4 5
2      2      6 7
Queue and Threshold Mapping for 1q4t (rx):
Queue Threshold CoS
-----
1      1      0 1
1      2      2
1      3      3 4 5
1      4      6 7
. . .

```

Example 3-4 Displaying MAC Information

```
cat6k-access> (enable) sh mac 5/1
```

Port	Rcv-Unicast	Rcv-Multicast	Rcv-Broadcast
5/1	267223	37	4

Port	Xmit-Unicast	Xmit-Multicast	Xmit-Broadcast
5/1	28748894	5206	72

Port	Rcv-Octet	Xmit-Octet
5/1	17178128	1840430081

"Out-Discards" are packets dropped due to congestion in the tx interface buffers

MAC	Dely-Exced	MTU-Exced	In-Discard	Out-Discard
5/1	0	0	0	262140

Example 3-5 Displaying QoS Summary Statistics

```
cat6k-access> (enable) sh qos stat 13
VoIP Control packets that have been re-written with CoS=3/DSCP=26 (AF31)
Packets dropped due to policing: 0
IP packets with ToS changed: 1885
IP packets with CoS changed: 781
Non-IP packets with CoS changed: 0
```

Example 3-6 Displaying QoS Detailed Statistics

```
cat6k-access> (enable) sh qos stat 5/1
All packets dropped are in the 1st drop threshold of queue #1
Tx port type of port 5/1 : 2q2t
Q # Threshold #:Packets dropped
---
1 1:393210 pkts, 2:0 pkts
2 1:0 pkts, 2:0 pkts
Rx port type of port 5/1 : 1q4t
Q # Threshold #:Packets dropped
---
1 1:0 pkts, 2:0 pkts, 3:0 pkts, 4:0 pkts
```

Uplink Interface to the Distribution Switch

Once you have configured all the access port queuing, you must also configure the uplink interfaces to the distribution/core switch. This involves enabling trust for Ethernet frames coming into the trunk port (1/1 in this example), manipulating the CoS-to-queue mapping entrance criteria, and mapping the CoS and IP Precedence values to the appropriate DSCP value. The procedure for doing this is outlined in the following sections.

MLS and Catalyst QoS Configuration

If the IP phones are in a different VLAN than Cisco CallManager, additional configuration is required. Any time a packet is sent to the Multilayer Switch Feature Card (MSFC) for Layer 3 switching, the CoS is set to 0. Because most configurations have the MSFC located in the distribution layer switch, the access layer switch must trust all DSCP tagging on the uplink trunk from the distribution layer. This enables the DSCP marking to be retained and used for DSCP-to-CoS Layer 3 classification in the wiring closet switch. Use **trust-cos** for Layer 2 uplinks and **trust-dscp** for Layer 3 uplinks; for example:

```
cat6k-access> (enable) set port qos 1/1 trust trust-dscp
```

Catalyst 6000 Transmit Queue Configuration

All VoIP (CoS=5) traffic will be placed into the egress interface Priority Queue on 1p2q2t interfaces and Queue 2 on 2q2t interfaces as soon as you enable QoS. However, you must perform the additional step of configuring the Catalyst 6000 CoS queue admission rules to ensure that CoS=3 (VoIP control) traffic is placed into the second queue. Use the following commands to perform this configuration:

```
cat6k-access> (enable) set qos map 1p2q2t tx 2 1 cos 3
cat6k-access> (enable) set qos map 2q2t tx 2 1 cos 3
```

Catalyst 6000 CoS/ToS-to-DSCP Mapping Configuration

Cisco follows the Internet Engineering Task Force (IETF) recommendations for setting the DSCP classification values for both the VoIP control plane traffic and VoIP bearer or media plane traffic. The recommended settings are DSCP=AF31 for VoIP control plane and DSCP=EF for VoIP bearer plane. To map the Layer 2 CoS and Layer 3 IP precedence settings correctly to these DSCP values, you must modify the default CoS/ToS-to-DSCP mappings as follows:

```
cat6k-distrib> (enable) set qos cos-dscp-map 0 8 16 26 32 46 48 56
cat6k-distrib> (enable) set qos ipprec-dscp-map 0 8 16 26 32 46 48 56
```

Verifying CoS/ToS-to-DSCP Mapping

To verify that the CoS and ToS settings are mapped correctly to the DSCP values, use the following two commands (shown with their associated output):

```
cat6k-distrib> (enable) sh qos map run cos-dscp-map
CoS - DSCP map:
CoS    DSCP
---    ----
0      0
1      8
2      16
3      26 -> 26 = AF31
4      32
5      46 -> 46 = EF
6 48
7 56

cat6k-distrib> (enable) sh qos map run ipprec-dscp-map
IP-Precedence - DSCP map:
IP-Prec  DSCP
-----  ----
0        0
1        8
2        16
3        26 -> 26 = AF31
4        32
5        46 -> 46 = EF
6        48
7        56
```

Catalyst 4000 Access Layer

Another popular campus configuration for Cisco AVVID networks uses Catalyst 2948G, 2980G, and 4000 series switches in the wiring closets. There are several compelling reasons for this:

- The Catalyst 4006 can provide in-line power to the IP phones.
- The Catalyst 4000 offers a very low price per port.
- These switches provide extremely scalable, high-speed switching.

Starting with CatOS Release 5.2, the Catalyst 4000 lines support dual-transmit queues on every interface. Admission to the queues is based on Layer 2 CoS markings and is configurable in 802.1p User Priority pairs.

Catalyst 4000 Port Scheduling and Queuing Schemes

This section presents several suggested configurations for port scheduling and queuing on the Catalyst 4000.

Receive Interface

The recommended configuration for the receive interface is

- FIFO

One standard FIFO (First-In, First-Out) queue.

Transmit Interface

The recommended configuration for the transmit interface is

- 2Q1T

Two standard queues with a single threshold. Scheduling is done on a Round-Robin (RR) basis. Admission to the queues is based on 802.1p CoS value and is user configurable in pairs. If you enable QoS but do not modify the CoS-to-transmit queue mappings, switch performance could be affected because all traffic is assigned to queue 1.



Note Once QoS is enabled on the Catalyst 4000, you must change the CoS mappings to utilize the newly created queue.

The default queue admission criteria for the Catalyst 4000 are

Queue #	Queue Admission CoS Value
1	0-7
2	Broadcast, Multicast, and Unknown Traffic

Figure 3-7 shows a general model for the Catalyst 4000 QoS configurations discussed in this guide.

Figure 3-7 General Model for Catalyst 4000 QoS Configurations

Catalyst 4000 Switch-Wide QoS

By default, only one queue is enabled on the Catalyst 4000 line of switches. Use the **set qos map** commands to enable the use of the second queue in CatOS Release 5.5.1. VoIP Control (CoS=3) frames should be placed into the second queue in the Catalyst 4000. These maps must be configured in pairs of CoS values because the Catalyst 4000 examines only the first two CoS bits; for example:

```
cat4k> (enable) set qos enable
cat4k> (enable) set qos map 2qlt 1 1 cos 0-1
cat4k> (enable) set qos map 2qlt 2 1 cos 2-3
cat4k> (enable) set qos map 2qlt 2 1 cos 4-5
cat4k> (enable) set qos map 2qlt 2 1 cos 6-7
```

Verifying Catalyst 4000 Queue Admission Configuration

To verify the queue admission configuration on the Catalyst 4000, use the following command (shown with its associated output):

```
cat4k> (enable) show qos info runtime
Run time setting of QoS:
QoS is enabled
All ports have 2 transmit queues with 1 drop thresholds (2qlt).
Default CoS = 0
Queue and Threshold Mapping:
Queue Threshold CoS
-----
1      1      0 1
2      1      2 3 4 5 6 7
```

IP Phone Port Queuing

In CatOS Release 5.5.1, the Catalyst 4000 line does not offer any advanced IP phone queuing features. Because of this, the Catalyst 4000 depends on the default CoS marking and enforcement on the IP phone. For more details, see Chapter 2, “Connecting IP Phones.”

Uplink Interface to the Distribution Switch

No special queuing or scheduling commands need to be configured on the Catalyst 4000 side of the link (from the access layer Catalyst 4000 to the distribution layer Catalyst 6000) because queuing is automatically enabled once QoS has been enabled and classification and queue admission have been configured.

You can perform additional uplink configuration if you are using the Catalyst 4000 with the Layer 3 engine (the WS-X4232, which enables IP, IPX, and Multicast routing for the switch). The Layer 3 engine enables the Catalyst 4000 to support four transmit queues based on IP precedence for

entrance criteria on the two-gigabit uplinks. The four queues are scheduled using a user-configurable WRR algorithm. In this case, the transmit interface configuration is as follows:

- 4Q1T

Four standard queues with a single threshold. Scheduling is done on a Round-Robin (RR) basis. Admission to the queues is based on 802.1p CoS value and is user configurable in pairs.



Note Once QoS is enabled on the Catalyst, you must change CoS mappings to utilize the newly created queue. Note that the Layer 3 queue numbering is the reverse of the Layer 2 numbering.

The default Layer 3 1000-Mbps uplink queue admission criteria for the Catalyst 4000 are as follows:

Queue #	Queue Admission IP Precedence Value
1	6-7
2	4-5
3	2-3
4	0-1

Catalyst 3500 Access Layer

The Cisco AVVID features in the Catalyst 2900 and Catalyst 3500 series, with a minimum Cisco IOS release of 12.0(5)XU, allow interaction with the IP phones for extending the CoS marking rules. In addition, the Catalyst 2900 XL and 3500 XL switches can classify untagged packets at the ingress ports by setting a default CoS priority for each port. However, these switches (except for the 3548 XL) cannot reclassify any tagged packets, and they honor only the 802.1p priority and place the packets in the appropriate transmit queue. All Catalyst 3500 switches and all Catalyst 2900 XL switches with 8-MB DRAM support these QoS features. The Catalyst 2900 XL with 4-MB DRAM does not support QoS features.

Catalyst 3500 Port Scheduling and Queuing Schemes

This section presents several suggested configurations for port scheduling and queuing on the Catalyst 3500.

Receive Interface

The recommended configuration for the receive interface is

- 1Q-FIFO

One standard FIFO (First-In, First-Out) queue.

Transmit Interface 10/100 Ports

The recommended configuration for the transmit interface for 10/100 ports is

- 2Q1T

Two standard queues with a single drop threshold. Scheduling is done on a priority-scheduling basis. Admission to the queues is based on 802.1p CoS or port-priority CoS value and is *not* user configurable.

The queue admission criteria for the Catalyst 3500 are as follows:

Queue #	Queue Admission CoS Value
1	0-3
2	4-7

Transmit Interface Gigabit Ethernet Ports

The recommended configuration for the transmit interface for gigabit Ethernet ports is

- 8Q-FIFO

Eight standard queues with a single drop threshold. Currently, only two queues are used. Scheduling is done on a priority-scheduling basis. Admission to the queues is based on 802.1p or port-priority CoS values and is *not* user configurable.

The gigabit Ethernet queue admission criteria are as follows:

Queue #	Queue Admission CoS Value
1	0-3
2	4-7
3-8	Not Used

Figure 3-8 shows a general model for the Catalyst 3500 QoS configurations discussed in this guide.

Figure 3-8 General Model for Catalyst 3500 QoS Configurations

IP Phone Port Queuing

If you use a single cable to install an IP phone, the access port is configured to trust the IP phone and not the attached PC. The port is also configured to use multiple transmit queues on all interfaces.

The commands to configure IP phone port queuing are

```
interface FastEthernet0/1
  switchport trunk encapsulation dot1q
  switchport trunk native vlan 12
  switchport mode trunk
  switchport voice vlan 112
  switchport priority extend cos 0
  spanning-tree portfast
```

Uplink Interface to the Distribution Switch

The recommended design for wiring closet configurations of Catalyst 3500 XL series switches is a star topology with a Catalyst 3524 PWR XL connected to a Catalyst 3508, which has dual uplinks to the distribution layer Catalyst 6000 switches. These uplinks are gigabit Ethernet links load balancing VLANs across the uplinks and configured with UplinkFast for fast Layer 2 convergence.

**Note**

A Catalyst 3500 series GigaStack configuration cannot provide guaranteed voice QoS because it is essentially a shared media access model.

The commands for this configuration are

```
interface GigabitEthernet0/1
  switchport trunk encapsulation dot1q
  switchport mode trunk
```

Catalyst 6000 Distribution Layer

After you configure the access switch and attach it to the distribution layer, you must set up Quality of Service (QoS) on the distribution switches. This requires the following changes to the configuration of the distribution switches:

- Configure VoIP control traffic transmit queuing.
- Configure the distribution layer with a Layer 3 access switch:
 - Enable trust ToS and DSCP from the access layer.
 - Configure ToS-to-DSCP mappings.
- Configure the distribution layer with a Layer 2 access switch:
 - Enable trust CoS and DSCP from the access layer.
 - Configure CoS-to-DSCP mappings.
 - Configure Layer 3 access lists for VoIP control traffic classification.
- Configure the connection to the Cisco 7200 WAN router.

Figure 3-9 shows a general model for these Catalyst 6000 distribution layer configurations, which are discussed in the following sections.

Figure 3-9 General Model for Catalyst 6000 Distribution Layer Configurations

Configuring Catalyst 6000 Distribution Layer VoIP Control Traffic Transmit Queue

As soon as QoS is enabled, all VoIP (CoS=5) traffic is placed into the egress interface Priority Queue on 1p2q2t interfaces and into Queue #2 on 2q2t interfaces (for all versions "1" of the 10/100 line cards). You must also perform an additional step of configuring the Catalyst 6000 CoS queue admission rules to ensure that CoS=3 traffic flows (VoIP control traffic) are placed into the second queue.

The commands for performing this configuration are

```
cat6k-distrib> (enable) set qos map 1p2q2t tx queue 2 1 cos 3
cat6k-distrib> (enable) set qos map 2q2t tx queue 2 1 cos 3
```

Catalyst 6000 Distribution Layer Configuration with a Catalyst 6000-PFC Access Layer

Once you have enabled QoS on the distribution layer switch and have modified the default queue admission, two configuration steps remain for completing the integration with an access layer switch that is enabled for Layer 3:

- Enable trust DSCP from the access layer.
- Configure ToS-to-DSCP mappings.

Trust DSCP from the Layer 3 Access Switch

Enable trust for DSCP values from adjacent Layer 3 access switches. Use port-base QoS on the trunking port and use **trust-dscp** instead of **trust-cos**. This is because **trust-cos** overwrites the Layer 3 DSCP value with the mapped CoS, and there is no need to do this since classification is done at the access layer.

The commands for this configuration are

```
cat6k-distrib> (enable) set port qos 1/1 port-based
cat6k-distrib> (enable) set port qos 1/1 trust trust-dscp
```

Catalyst 6000 ToS-to-DSCP Mapping Configuration

Cisco follows the Internet Engineering Task Force (IETF) recommendations for setting the DSCP classification values for both the VoIP control plane traffic and VoIP bearer or media plane traffic. The recommended settings are DSCP=AF31 for VoIP control plane and DSCP=EF for VoIP bearer plane. To map the Layer 3 IP precedence settings correctly to these DSCP values, you must modify the default ToS-to-DSCP mappings as follows:

```
cat6k-distrib> (enable) set qos ipprec-dscp-map 0 8 16 26 32 46 48 56
```

```
cat6k-distrib> (enable) sh qos map run ipprec-dscp-map
```

```
IP-Precedence - DSCP map:
IP-Prec      DSCP
-----
0            0
1            8
2           16
3           26 -> 26 = AF31
4           32
5           46 -> 46 = EF
6           48
7           56
```

Catalyst 6000 Distribution Layer Configuration with an Access Switch Enabled for Layer 2 Only

Once you have enabled QoS on the distribution layer switch and have modified the default queue admission, you must perform three additional configuration steps to complete the integration with a Layer 2 access switch:

- Enable trust CoS from the access layer.
- Configure CoS-to-DSCP mappings.
- Configure Layer 3 access lists for VoIP control traffic classification. (See the “Marking Control and Management Traffic” section on page 3-6.)

Trust CoS from the Layer 2 Access Switch

Enable trust for CoS values from adjacent Layer 2 access switches. Use **vlan-based** QoS on the trunking port and use **trust-cos** instead of **trust-dscp**. This configuration is used when the access layer switch is only a Layer 2 device doing CoS classification.

The command for this configuration is

```
cat6k-distrib> (enable) set port qos 1/2,3/2 vlan-based
cat6k-distrib> (enable) set port qos 1/2,3/2 trust trust-cos
```

Catalyst 6000 CoS-to-DSCP Mapping Configuration

Cisco follows the IETF recommendations for setting the DSCP classification values for both the VoIP control plane traffic and VoIP bearer or media plane traffic. The recommended settings are DSCP=AF31 for VoIP control plane and DSCP=EF for VoIP bearer plane. To map the Layer 2 settings correctly to these DSCP values, you must modify the default CoS-to-DSCP mappings as follows:

```
cat6k-distrib> (enable) set qos cos-dscp-map 0 8 16 26 32 46 48 56
```

```
cat6k-distrib> (enable) sh qos map run cos-dscp-map
```

```
CoS - DSCP map:
CoS   DSCP
---   ----
0     0
1     8
2    16
3    26 -> 26 = AF31
4    32
5    46 -> 46 = EF
6    48
```

Configuring Layer 3 Access Lists for VoIP Control Traffic Classification

To configure Layer 3 access lists for VoIP Control traffic classification, use the following commands (shown with their associated output). Also use the ACL_IP-PHONES access list from the “Marking Control and Management Traffic” section on page 3-6.

```
cat6k-distrib> (enable) set port qos 1/2,3/2 vlan-based
cat6k-distrib> (enable) set qos acl map ACL_IP-PHONES 111
```

```
cat6k-distrib> (enable) sh qos acl map run ACL_IP-PHONES
ACL name                               Type Vlans
```

```
-----
ACL_IP-PHONES                           IP 110,111,112
```

```
ACL name                               Type Ports
```

```
-----
ACL_IP-PHONES                           IP
```

```
cat6k-distrib> (enable) sh qos acl info run ACL_IP-PHONES
```

```
set qos acl IP ACL_IP-PHONES
```

```
-----
1. dscp 26 tcp any any range 2000 2002
2. dscp 26 tcp any any eq 1720
3. dscp 26 tcp any any range 11000 11999
4. dscp 26 udp any any eq 2427
5. trust-cos any
```



Note

Beginning with Release 3.0(5), Cisco CallManager includes the ability to configure the CoS and ToS values for all VoIP control and management traffic from Cisco CallManager, the IP phones, and the Skinny Protocol gateways (this does not include the AT and AS model analog gateways). With this user-configurable classification, network element access lists are no longer required for marking Skinny Protocol VoIP control traffic. H.323 and Media Gateway Control Protocol (MGCP) traffic still require external, network element marking for several more months.

Configuring the Connection to the Cisco 7200 WAN Router

Use the following commands to configure the connection to the Cisco 7200 WAN router.



Note

Current 10/100 version "1" line cards must still have **trust-ipprec** enabled even though the parser returns an error

```
cat6k-distrib> (enable) set port qos 9/1 port-based
cat6k-distrib> (enable) set port qos 9/1 trust trust-ipprec
```

```
cat6k-distrib> (enable) set qos acl ip ACL_TRUST-WAN trust-ipprec any
cat6k-distrib> (enable) commit qos acl ACL_TRUST-WAN
cat6k-distrib> (enable) set qos acl map ACL_TRUST-WAN 9/1
```

```
cat6k-distrib> (enable) sh port qos 9/1
QoS is enabled for the switch.
QoS policy source for the switch set to local.
```

Port	Interface config	Type	Interface runtime	Type	Policy config	Source	Policy runtime	Source
9/1	port-based		port-based			COPS		local

Port	TxPort	Type	RxPort	Type	Trust config	Type	Trust runtime	Def	CoS	Def	CoS
									config		runtime
9/1		2q2t	1q4t		trust-ipprec		trust-ipprec		0		0

Port	Ext-Trust	Ext-Cos
9/1	untrusted	0

(*)Runtime trust type set to untrusted.

```
Config:
Port  ACL name                                Type
-----
9/1   ACL_TRUST-WAN                            IP
```

```
Runtime:
Port  ACL name                                Type
-----
9/1   ACL_TRUST-WAN                            IP
```

Catalyst 6000 Distribution/Core Running Native IOS

After you configure the access switch and attach it to the distribution layer, you must set up Quality of Service (QoS) on the distribution switches. This requires the following changes to the configuration of the distribution switches:

- Configure QoS.
- Configure VoIP control traffic transmit queuing.
- Configure the distribution layer with a Layer 3 access switch:
 - Enable trust ToS and DSCP from the access layer.
 - Configure ToS-to-DSCP mappings.
- Configure the distribution layer with a Layer 2 access switch:
 - Enable trust CoS and DSCP from the access layer.
 - Configure CoS-to-DSCP mappings.
 - Configure the QoS policies and Layer 3 access lists for VoIP control traffic classification.

Figure 3-10 shows a general model for configurations running Native Cisco IOS on Catalyst 6000 distribution layer switches. The configuration details are discussed in the sections that follow.

Figure 3-10 General Model for Native Cisco IOS Running on Catalyst 6000 Distribution Layer Switches

Configuring QoS on the Native Cisco IOS Catalyst 6000

To enable QoS on the Catalyst 6000 with Native Cisco IOS, use the following command:

```
mls qos
```

Configuring Transmit Queue Admission for VoIP Control Traffic

As soon as QoS is enabled, all VoIP (CoS=5) traffic is placed into the egress interface Priority Queue on 1p2q2t interfaces and into Queue #2 on 2q2t interfaces (for all versions "1" of the 10/100 line cards). You must also perform an additional step of configuring the Catalyst 6000 CoS queue admission rules to ensure that CoS=3 traffic flows (VoIP control traffic) are placed into the second queue.

The commands for performing this configuration are

```
int range gigabitEthernet 1/1 - 2
  wrp-queue cos-map 1 2 2
  wrp-queue cos-map 2 1 3 4
```

Catalyst 6000 Native Cisco IOS Distribution Layer Configuration with a Catalyst 6000-PFC Access Layer

Once you have enabled QoS on the Native Cisco IOS distribution layer switch and have modified the default queue admission, two configuration steps remain for completing the integration with an access layer switch that is enabled for Layer 3:

- Enable trust DSCP from the access layer.
- Configure ToS-to-DSCP mappings.

Trust DSCP from the Layer 3 Access Switch

Enable trust for DSCP values from adjacent Layer 3 access switches. Use port-base QoS on the trunking port (port-based QoS is enabled by default when **mls qos** is configured), and use **mls qos trust dscp** instead of the CatOS **trust-dscp**.



Note

Classification has already been established at the access layer in this model.

The commands for this configuration are

```
interface GigabitEthernet2/1
  description trunk port to PFC enabled cat6k-access
  no ip address
  wrr-queue cos-map 1 2 2
  wrr-queue cos-map 2 1 3 4
  mls qos trust dscp
  switchport
  switchport trunk encapsulation dot1q
  switchport mode trunk
```

Native Cisco IOS ToS-to-DSCP Mapping Configuration for Layer 3 Access Switches

Cisco follows the IETF recommendations for setting the DSCP classification values for both the VoIP control plane traffic and VoIP bearer or media plane traffic. The recommended settings are DSCP=AF31 for VoIP control plane and DSCP=EF for VoIP bearer plane. To map the Layer 3 IP precedence settings correctly to these DSCP values, you must modify the default ToS-to-DSCP mappings.



Note

The Catalyst 6000 numerical values of 26 and 46 correlate to DSCP=AF31 and DSCP=EF, respectively. This is done in global configuration mode.

The command for this configuration is

```
mls qos map ip-prec-dscp 0 8 16 26 32 46 56 0
```

Catalyst 6000 Native Cisco IOS Distribution Layer Configuration with an Access Switch Enabled for Layer 2 Only

Once you have enabled QoS on the distribution layer switch and have modified the default queue admission, you must perform three additional configuration steps to complete the integration with a Layer 2 access switch:

- Enable trust CoS from the access layer.
- Configure CoS-to-DSCP mappings.
- Configure Layer 3 access lists for VoIP control traffic classification. (See the “Marking Control and Management Traffic” section on page 3-6.)

Trust CoS from the Layer 2 Access Switch

Enable trust for CoS values from adjacent Layer 2 access switches. Use port-base QoS on the trunking port, and use the Native IOS command **mls qos trust cos** instead of CatOS **trust-cos**. This configuration is used when the access layer switch is only a Layer 2 device doing CoS classification.

The commands for this configuration are

```
interface GigabitEthernet2/2
  description trunk port to layer 2-only cat4k
  no ip address
  wrr-queue cos-map 1 2 2
  wrr-queue cos-map 2 1 3 4
  mls qos vlan-based
  mls qos trust cos
  switchport
  switchport trunk encapsulation dot1q
  switchport mode trunk
!
interface GigabitEthernet3/1
  description trunk port to layer 2-only 3500
  no ip address
  wrr-queue cos-map 1 2 2
  wrr-queue cos-map 2 1 3 4
  mls qos vlan-based
  mls qos trust cos
  switchport
  switchport trunk encapsulation dot1q
  switchport mode trunk
```

Native IOS CoS-to-DSCP Mapping Configuration for Layer 2 Access Switches

Cisco follows the IETF recommendations for setting the DSCP classification values for both the VoIP control plane traffic and VoIP bearer or media plane traffic. The recommended settings are DSCP=AF31 for VoIP control plane and DSCP=EF for VoIP bearer plane. To map the Layer 2 settings correctly to these DSCP values, you must modify the default CoS-to-DSCP mappings.

**Note**

The Catalyst 6000 numerical values of 26 and 46 correlate to DSCP=AF31 and DSCP=EF, respectively. This is done in global configuration mode.

The command for this configuration is

```
mls qos map cos-dscp 0 8 16 26 32 46 56 0
```

Configure the QoS Policies and Layer 3 Access Lists for VoIP Control Traffic Classification

The QoS configuration for the Native Cisco IOS Catalyst 6000 is very similar to the WAN router Cisco IOS configurations, with the exception of using policing for marking traffic flows and applying service policies to VLAN interfaces. The physical gigabit Ethernet uplink ports are configured to use VLAN-based QoS with the **mls qos vlan-based** Native Cisco IOS interface commands. Finally, the **service-policy** is applied to all VLAN traffic *inbound* on the uplink.

In the following example, three classes are defined: one for the VoIP media stream, one for the control traffic, and the last for all other traffic. Traffic is filtered for these classes based on Layer 3 or 4 source and destination IP addresses and ports. Each of these classes is referenced in the **Voice-QoS** policy map. In the **policy-map** statements, a policing function is used to classify all traffic that meets the entrance criteria matched with the **class-map** access lists.

**Note**

The Catalyst 6000 Native Cisco IOS software does not support the **set ip dscp** commands. Instead, the policing algorithm is used for classifying traffic.

In this scenario, the policing code tags the traffic flows with DSCP values of AF31 for VoIP control traffic, EF for VoIP Media traffic, and 0 for all other packets. The size of the "8000" flows is low enough that any traffic will solicit tagging using the syntax **conform-action set-dscp-transmit 26**.

```
class-map match-all VoIP-Control
  match access-group 100
class-map match-all VoIP-RTP
  match access-group 101
class-map match-all Routine
  match access-group 102
!
!
policy-map Voice-QoS
  class VoIP-Control
    police 8000 8000 8000 conform-action set-dscp-transmit 26 exceed action transmit
  class VoIP-RTP
    police 8000 8000 8000 conform-action set-dscp-transmit 46 exceed-action transmit
  class Routine
    police 8000 8000 8000 conform-action set-dscp-transmit 0 exceed-action transmit
!
! access-list 100 looks for VoIP Control Traffic
access-list 100 permit tcp any any range 2000 2002
access-list 100 permit tcp any any eq 1720
access-list 100 permit tcp any any range 11000 11999
access-list 100 permit udp any any eq 2427
!
! access-list 101 looks for VoIP Bearer Traffic
access-list 101 permit udp any any range 16384 32767
!
! access-list 102 filters for routine traffic
access-list 102 permit ip any any
!
interface GigabitEthernet2/2
  description trunk port to layer 2-only cat4k
  no ip address
  wrr-queue cos-map 1 2 2
  wrr-queue cos-map 2 1 3 4
  ! inform the port that QoS will be VLAN-Based
  mls qos vlan-based
  switchport
  switchport trunk encapsulation dot1q
  switchport mode trunk
!
```

```
interface GigabitEthernet3/1
  description trunk port to layer 2-only 3500
  no ip address
  wrr-queue cos-map 1 2 2
  wrr-queue cos-map 2 1 3 4
  ! inform the port that QoS will be VLAN-Based
  mls qos vlan-based
  switchport
  switchport trunk encapsulation dot1q
  switchport mode trunk
!
interface Vlan111
  description voice vlan on cat4k
  ip address 10.1.111.77 255.255.255.0
  ip helper-address 10.1.10.10
  no ip redirects
  ! apply the QoS policy as an inbound policy
  service-policy input Voice-QoS
  standby 111 ip 10.1.111.1
!
interface Vlan112
  description voice vlan on 3500
  ip address 10.1.112.77 255.255.255.0
  ip helper-address 10.1.10.10
  no ip redirects
  ! apply the QoS policy as an inbound policy
  service-policy input Voice-QoS
  standby 112 ip 10.1.112.1

ios6k#sh mls qos
QoS is enabled globally
Microflow policing is enabled globally

QoS is vlan-based on the following interfaces:
Vl111 Vl112 Gi2/2 Gi3/1 Gi3/2 Gi3/3
Gi3/4 Gi3/5 Gi3/6 Gi3/7 Gi3/8 Gi4/1 Gi4/2 Gi4/3 Gi4/4 Gi4/5
Gi4/6 Gi4/7 Gi4/8 Fa9/1 Fa9/2 Fa9/3 Fa9/4 Fa9/5 Fa9/6 Fa9/7
Fa9/8 Fa9/9 Fa9/10 Fa9/11 Fa9/12 Fa9/13 Fa9/14 Fa9/15 Fa9/16 Fa9/17
Fa9/18 Fa9/19 Fa9/20 Fa9/21 Fa9/22 Fa9/23 Fa9/24 Fa9/25 Fa9/26 Fa9/27
Fa9/28 Fa9/29 Fa9/30 Fa9/31 Fa9/32 Fa9/33 Fa9/34 Fa9/35 Fa9/36 Fa9/37
Fa9/38 Fa9/39 Fa9/40 Fa9/41 Fa9/42 Fa9/43 Fa9/44 Fa9/45 Fa9/46 Fa9/47
Fa9/48
```

QoS global counters:

```
Total packets: 16750372458300
Packets dropped by policing: 55930847232
IP packets with TOS changed by policing: 16750372458300
IP packets with COS changed by policing: 55945330688
Non-IP packets with COS changed by policing: 16750372458300
```

Summary

As described in this chapter, the following general guidelines and recommendations apply when configuring a Cisco AVVID network in a campus environment:

- Multiple queues are required on all interfaces to guarantee voice quality.
- To enable fast convergence, use UplinkFast in wiring closet switches that have multiple egress queues, such as the Catalyst 2900 XL, 3500, 2948G, 2980G, 4000, and 6000 switches.
- Set all Cisco AVVID control and management traffic to maximum CoS and ToS values of 3.
- Never allow PC applications to send traffic at a CoS or ToS value of 4-7.
- Distribution layer switches must have the ability to map Layer 3 ToS to Layer 2 CoS values.